A discrete-time model of human motor learning

Brandon ROHRER[1] and Steven HULET[1]

[1] *Intelligent Systems and Robotics Center,*

Sandia National Laboratories, USA. brrohre@sandia.gov

Abstract. Current models of human motor learning and control typically employ continuous (or near continuous) movement commands and sensory information. However, research suggests that voluntary motor commands are issued in discrete-time submovements. There is also reasonable support for the hypothesis that human sensory experience is episodic as well. These facts have motivated the development of a learning model that employs discrete-time sensory and motor control events. We present this model together with the results of initial simulation of robot control. The results show that the learning that takes place is adaptive and is robust to a variety of conditions that many traditional controllers are not capable of handling, including random errors in the actuators and sensors, random transmission time delays, hard nonlinearities, time varying system behavior, and unknown structure of system dynamics.

Full Research Paper

Conference topics: Neuroscience, Handwriting Generation

Keywords: motor control, motor learning, adaptive control, event-based control, submovements

1. Introduction

When mathematically modeling human motor learning and control, it is common to make a number of assumptions. Sensory information and control information are usually considered to be continuous in time. Perception and movements are often expressed in terms of world coordinates. In many cases, velocity, acceleration, and higher derivatives of position are explicitly represented in the motion planner. And finally, kinematic states are assumed to be sensed at high resolution. While models based on these assumptions can describe some aspects of human movement, none of these assumptions has been proven. In addition, these models are typically used only to model and predict a limited class of movement (e.g., reaching movements). In this paper, we propose an alternative motor learning model. This model employs as working assumptions that both motor commands and sensory information are passed in an episodic fashion, quantized in time.

Evidence for discrete time motor commands, also known as submovements, is widespread and accounts for a large number of disparate phenomena in motor behavior. Observations of slow finger movements (Vallbo & Wessberg, 1993), eye saccades (Collewijn et al., 1988), tracing constant curvature paths (Abend et al., 1982), cyclical movements (Woodworth, 1899; Crossman & Goodeve, 1983; Doeringer, 1999), infant reaching movements (Hofsten, 1991), ballistic movements (Morasso & Mussa-Ivaldi, 1982), movements of recovering stroke patients (Krebs et al., 1999; Rohrer et al., 2004), and movements requiring high accuracy (Milner, 1992) are all consistent with a theory of submovements. The discrete-time nature of movement is evident not only in movement kinematics, but also in the electromyograph (EMG) signals of agonist and antagonist muscles (Vallbo & Wessberg, 1993).

Evidence for the discrete nature of sensory experiences is more subtle. The concept was originally proposed by William James (1890) and more recently in Stroud (1956). A number of experiments support a theory of discrete sensory experiences. A striking phenomenon that suggests discrete sensory experience is the wagon wheel illusion under steady light. Due to the rapid series of photographs of which movies are composed, it is commonly observed that a spoked wagon wheel appears to rotate slowly backward while rolling rapidly forward. Interestingly, the same effect can also be observed under steady light (Purves et al., 1996), suggesting a periodic sampling mechanism in human vision. In another experiment, two lights that blinked with a slight delay were occasionally perceived to flash simultaneously (Wertheimer, 1912), an occurrence that was suggested to be a function of the phase relationship with alpha (8–12 Hz) cortical rhythms (Gho & Varela, 1988). Other observations that suggest discrete sensory experiences are the sharp dependence of perceived causality on delay times and periodicities in reaction times (VanRullen & Koch, 2003). A more in-depth review of the case for discrete perception is made in Koch (2004).

2. An event based motor learning agent

Because motor commands and sensory signals in the models occur at discrete intervals, they can also be termed motor and sensory "events," and any model employing them can be described as "event-based." For the sake of implementation, we developed an instance of an event-based learning agent in which several additional limitations are imposed. (1) Motor commands and sensory information are not only quantized in time, but are also coarsely discretized in magnitude. (2) Motor and sensory events of different magnitudes are considered categorically unrelated. That is, extrapolation and interpolation do not occur explicitly. (3) Motor and sensory events are serially registered in the motor learning agent.

Learning occurs by repeated observation of sensory and control events. The block diagram in figure 1 describes the process in detail. During learning (as in an infant), the motor control system issues random commands and observes the resulting sensory events. As patterns are observed repeatedly, they are recorded and extended. This growing library of patterns constitutes the motor controller's "experience base."

3. Simulation

Consider an implementation of the learning model outlined above in a simulated rotary pointer robot. Possible sensory events for the pointer robot consist of position sensing in 10° bins, resulting in 36 distinct states (for convenience, numbered 1 through 36). Possible command events for the pointer are 10° rotation clockwise (R) and and 10° rotation counterclockwise (L).

An example of how the learning model operates shows the simplicity of the approach. One sample excerpt of an event history resulting from random movements might consist of 3R4R5L4R5L4L3L2. The learning agent would break the event history into short patterns, 3R4, 4R5, 5L4, etc. When these patterns are encountered again in subsequent excerpts of the event history, they will be extended, producing patterns such as 3R4R5, 4R5L4, 5L4R5R6.

A simulation of the learning agent using the pointer robot system was implemented in C++. Six conditions were simulated:

Simple system. Measurement states 1–36 and command events R and L as described previously.

Hard stop. Same as the simple system, but with a "hard stop" inserted at 0°, prohibiting continuous rotational movement.

Sensory state scramble. Same as the simple system, but after 5000 trials, the numerical labels for sensory states are renamed 1–36 in random order, making all prior learning inapplicable and misleading.

Command reversal. Same as the simple system, but after 5000 trials the commands "reverse"; an R command produces *counterclockwise* motion and an L command produces *clockwise* motion.

Random error. Same as the simple system, but with up to 5° of random error added to each command event, resulting in movements of between 5° and 15°. With measurement resolution limited to 10°, the error will express itself as measurement states being either skipped or unchanged when a command is issued.

Random delays. Same as the simple system, but each command event has a 50% chance of being delayed and executed at the instant the *next* command is issued. As a result, when a command is issued, zero, one, or two command events may actually take place.

In each case, the learning agent generated random command events and attempted to predict the results before executing the command. Predictions were generated by searching through previously observed patterns for instances containing a portion of the current event history. Patterns that matched a longer portion of the event history were favored heavily. Patterns that were observed recently or that had been observed many times were also favored. Once a pattern was selected a prediction was obtained by reading "what happened next" when the situation had been encountered previously. In each condition, the learning agent began with a clean slate;

that is, there were no previously observed experiences upon which to build. As a result, lack of prior experience made it impossible for the agent to offer a prediction in some cases. These were counted as unsuccessful predictions.

4. Results

During simulation of each of the six conditions, the learning agent generated a database of patterns. Typical patterns observed were 16 R 17 R 18 R 17 L 18 R 19 (observed 6 times), 25 R 26 R 27 L 26 L 25 R 26 (observed 8 times), and 22 R 23 L 22 R 23 L 22 L 21 (observed 11 times). In the case of the simple system, a total of 2155 repeated patterns were observed, occupying 599 kilobytes of memory. The longest patterns observed included five movement events, a limit imposed by the software, rather than by the inherent function of the learning agent. On average, patterns contained between 3 and 4 movement events. Simulating 10,000 trials for one condition took approximately two minutes. Given that the simple system was learned within the first 2000 trials, only 24 seconds were required to learn the system's dynamics completely.

The results of the simulations are shown in figure 2. As shown in the plot, the learning agent achieved 100% accuracy in the simple system after 2000 trials. The learning agent showed similar performance in the presence of a hard stop. In both these conditions, the performance of the system is deterministic, allowing correct predictions at every time step.

Scrambling the sensory state labels changed the system fundamentally, making the probability of encountering a previously observed pattern small. Learning essentially began from scratch, and the initial learning transient was repeated after scrambling. Reversing the directions of the commands resulted in a marked decrease in performance initially, but the agent recovered within 4000 trials after that and predicted the last 1000 trials perfectly.

The introduction of random noise into the movement amplitude made perfect prediction impossible. The noise amplitude exactly corresponded to the resolution of the position measurement, 10°. As a result, knowledge of the current position allowed prediction the subsequent position with an accuracy of only 50%. With a longer event history, it was possible to increase the accuracy, but only to a certain extent. The learning agent began with a prediction accuracy slightly higher than 50%, and gradually it increased to near 70%.

Random time delays introduced the possibility that zero, one, or two command events might be executed at once. With a 50% probability of delay, at any given time step there was a 25% chance that no command would be executed, a 25% chance that two commands would be executed simultaneously, a 25% chance that the previous command alone would be executed, and only a 25% chance that the current command alone would be executed. As a result, even once the behavior of this simple system is learned, only a 25% success rate can be expected with no knowledge of prior events. However, with a complete knowledge of prior events, it was possible to infer whether the prior command event had been executed, allowing a prediction accuracy of 50%. The learning agent began with prediction accuracy slightly higher than 25%, and that accuracy climbed to just over 45% after 10,000 trials.

5. Discussion

The computational requirements of the learning agent were modest. Although the system simulated was simple, the time required to learn its dynamics fully was less than 30 seconds and data storage requirements were almost negligible (¡1MB). However, we anticipate that both the learning time and the storage requirements will increase exponentially with the number of possible sensory and command events. Extending the learning agent to more complicated systems may require strategies to limit the number of possible events, or to introduce new sensory and motor events gradually.

It is worth noting that the prediction tasks demonstrated here are nontrivial. The five conditions contained instances of hard nonlinearities, dramatic time variance, large stochastic movement error, and nondeterministic time delays, any one of which can impose insurmountable challenges for certain learning algorithms. However, they are challenges that the human motor learning mechanism routinely faces and successfully overcomes without difficulty. Taken together, they constitute something of a proving ground for any model of motor learning purporting to describe that of a human.

6. References

Abend, W., Bizzi, E., & Morasso, P. (1982). Human arm trajectory formation. *Brain, 105*, 331-348.

Collewijn, H., Erkelens, C. J., & Steinman, R. M. (1988). Binocular coordination of human horizontal saccadic eye movements. *J Physiol, 404*, 157-182.

Crossman, E. R. F. W., & Goodeve, P. J. (1983). Feedback control of hand-movements and fitt's law. *Quarterly Journal of Experimental Psychology, 35A*, 251-278. (Paper presented at the meeting of the Experimental Psychology Society, Oxford, July 1963)

Doeringer, J. A. (1999). *An investigation into the discrete nature of human arm movements.* Unpublished doctoral dissertation, Massachusetts Institute of Technology.

Gho, M., & Varela, F. J. (1988). A quantitative assessment of the dependency of the visual temporal frame upon the cortical rhythm. *J Physiol. Paris, 83*, 95-101.

Hofsten, C. von. (1991). Structuring of early reaching movements: A longitudinal study. *Journal of Motor Behavior*, *23*(4), 280-292.

James, W. (1890). *The principles of psychology.* New York: Dover Publications.

Koch, C. (2004). *The quest for consciousness.* Englewood, Colorado: Roberts & Company Publishers.

Krebs, H. I., Aisen, M. L., Volpe, B. T., & Hogan, N. (1999). Quantization of continuous arm movements in humans with brain injury. *Proceedings of the National Academy of Sciences USA, 96*, 4645-9. (Neurobiology)

Milner, T. E. (1992). A model for the generation of movements requiring endpoint precision. *Neuroscience*, *49*, 365-374.

Morasso, P., & Mussa-Ivaldi, F. A. (1982). Trajectory formation and handwriting: A computational model. *Biological Cybernetics*, *45*, 131-142.

Purves, D., Paydarfar, J. A., & Andrews, T. J. (1996). The wagon wheel illusion in movies and reality. *Proceedings of the National Academy of Sciences USA*, *93*, 3693-3697.

Rohrer, B., Fasoli, S., Krebs, H. I., Volpe, B., Frontera, W. R., Stein, J., & Hogan, N. (2004). Submovements grow larger, fewer, and more blended during stroke recovery. *Motor Control*, *8*(4), 472-483.

Stroud, J. M. (1956). Information theory in psychology. In H. Quastler (Ed.), (p. 174-205). Glencoe IL: Free Press.

Vallbo, A. B., & Wessberg, J. (1993). Organization of motor output in slow finger movements in man. *Journal of Physiology*, *469*, 673-691.

VanRullen, R., & Koch, C. (2003). Is perception discrete or continuous? *Trends in Cognitive Sciences*, *7*, 207-213.

Wertheimer, M. (1912). Experimentelle studien über das sehen von bewegung. *Z. Psychologie*, *61*, 161-265.

Woodworth, R. S. (1899). The accuracy of voluntary movement. *Psychology Review Monogr Suppl*.

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

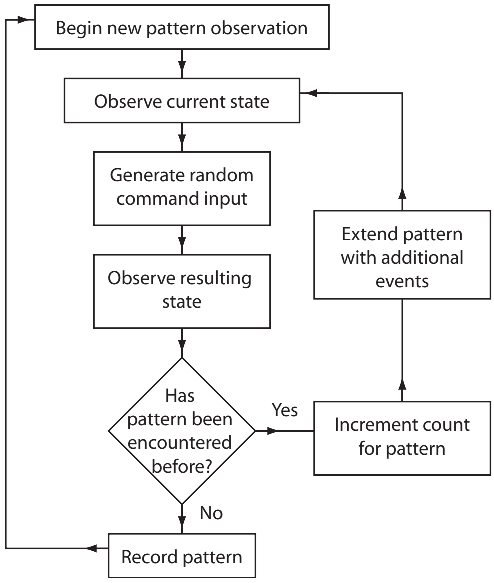

Figure 1. Learning agent operation. The learning agent identifies its dynamics and those of its environment by randomly generating patterns and noting repeated occurrences. It initially observes its current state, generates a random input, and then observes the resulting state. If it is a pattern that it has not encountered before, it records the pattern in memory and repeats the process. If the pattern has been previously observed, the agent notes the observation and then extends the pattern by generating another random input. In this way, patterns of increasing length are recorded as training progresses.

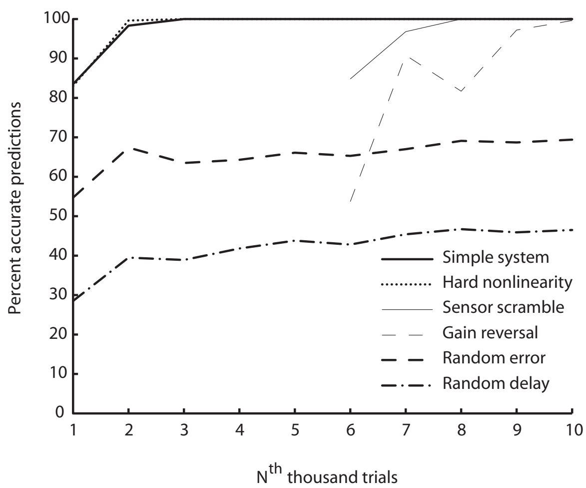

Figure 2. Simulation performance. Six different conditions were imposed: the simple system (bold solid line), the system with a hard stop (dotted line), the system with scrambled sensory state labels (fine solid line), the system with a command reversal after 5000 trials (fine dotted line), the system with randomness in the movement amplitude (dashed line), and the system with a random time delay (dash-dot line).